TactileLab: Efficient Shear-Sensitive Tactile Simulation for Dexterous Sim2Real Robotic Manipulation

Yijiong Lin, Nathan F. Lepora
School of Engineering Mathematics and Technology, University of Bristol
Bristol Robotics Laboratory, University of Bristol

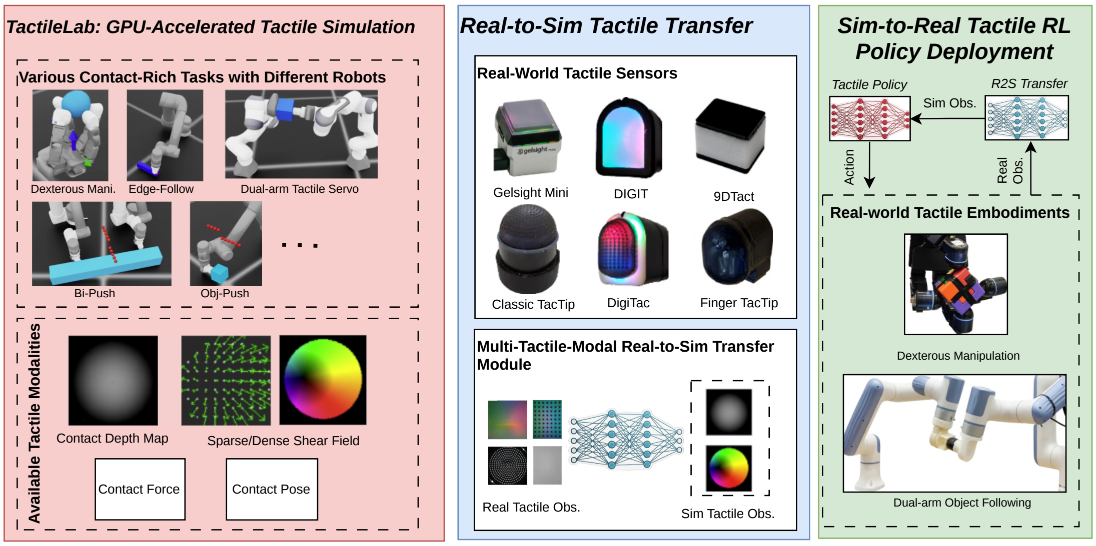

Figure 1. TactileLab is an end-to-end framework for efficient shear-sensitive tactile robot learning. We developed a GPU-accelerated tactile simulation based on IsaacLab for diverse contact-rich tasks and heterogeneous robots, tactile learning modules for policy and representation learning, a multi-modal real-to-sim tactile transfer pipeline for multiple real tactile sensors, and sim-to-real tactile RL policy deployment on real robotic systems. This unified design supports scalable training, efficient simulated tactile data collection, and transfer from simulation to real-world tactile manipulation.

I. INTRODUCTION

Tactile sensing provides contact-grounded feedback that is essential for contact-rich manipulation. This is especially important when a robot must maintain stable contact, regulate interaction forces, or respond to small changes at the contact interface. Vision-based tactile sensors are attractive because they provide high-resolution spatial observations and rich temporal measurements. However, scalable simulation methods still struggle to represent tangential interaction efficiently.

Most existing tactile simulators fall into two categories. Deformable simulators can model compliant contact with high fidelity [1], but are often too slow for large-scale reinforcement learning. Rigid-body simulators are fast and scalable, but usually provide only depth-like, force-like, or compact contact signals [2]–[6]. As a result, they are well suited for normal-contact reasoning but less effective for tasks in which shear, slip, or tangential motion matters. This limitation is important for manipulation problems such as in-contact trajectory following, insertion, and in-hand object manipulation under external perturbations. In these tasks, a depth-only observation cannot fully describe how contact evolves over time. Different contact states may produce similar instantaneous depth signals while leading to very different future outcomes, causing partial observability and making policy learning more difficult. Although recent tactile simulation frameworks such as TacSL can model shear-related tactile quantities [7], they still rely on simulator-specific assumptions and object geometry, which can limit generalization and real-world deployment.

To address this gap, we introduce **TactileLab**, an efficient tactile simulation and learning framework that augments rigid-body tactile sensing with dense tangential contact-motion cues. Rather than explicitly simulating elastomer deformation, TactileLab represents contact using complementary tactile modalities: normal-contact depth and simulated tactile flow. Contact depth captures local indentation and contact geometry, while tactile flow provides a dense in-plane motion field that describes tangential contact evolution. This representation is efficient to compute, compatible with large-scale reinforcement learning, and transferable to real vision-based tactile sensing through real-to-sim tactile translation.

II. METHODS

A. TactileLab Framework

TactileLab is an end-to-end framework for tactile robot learning built on IsaacLab [8]. It unifies efficient tactile simulation, contact-rich task design, tactile reinforcement learning, real-to-sim tactile transfer, and sim-to-real deployment within a single GPU-parallelized pipeline. The framework is designed to support heterogeneous robotic embodiments, including single-arm, dual-arm, gripper-based, and multi-finger platforms, while maintaining a common interface for tactile observation generation, policy learning, and deployment.

At the system level, TactileLab contains four main components. First, it provides an efficient tactile simulation stack with multiple tactile modalities, including low-dimensional contact information, contact-depth images, and simulated tactile-flow fields. Second, it offers a unified library of tactile robotic tasks spanning manipulation regimes from non-prehensile interaction to dexterous in-hand control. Third, it includes learning tools for tactile reinforcement learning, teacher–student distillation, and tactile data collection. Fourth, it supports real-to-sim tactile transfer and real-world deployment, enabling policies trained with simulated tactile observations to operate from real tactile inputs. Together, these components make TactileLab a scalable platform for studying contact-rich manipulation under a shared simulation-to-deployment workflow.

B. Shear-Sensitive Tactile Representation

The key representation in TactileLab augments contact depth with simulated tactile flow, a dense tangential motion field defined on the tactile image plane. Contact depth captures normal interaction, while tactile flow captures local in-plane motion at each spatial location. Together, they form a shear-sensitive tactile observation that encodes both contact geometry and tangential contact evolution.

This design is motivated by the limitation of depth-only tactile sensing in shear-dependent tasks. A depth map can indicate how deeply the object indents into the tactile surface, but it does not explicitly describe how the contact patch moves, drifts, or changes tangentially. As a result, different contact states may produce similar instantaneous depth observations while corresponding to different future dynamics. By incorporating tactile flow, the observation becomes more informative for tasks that require tracking relative motion, regulating shear, recovering from slip-like disturbances, or reasoning about evolving contact conditions.

Importantly, simulated tactile flow is defined as an efficient abstraction rather than a full deformable-material simulation. Instead of explicitly solving elastomer mechanics, the framework represents shear-relevant interaction through image-space tangential motion cues induced by contact-interface motion. This preserves the scalability of rigid-body simulation for reinforcement learning while providing a tactile modality that can be aligned with optical-flow-like signals in real camera-based tactile sensors.

C. Contact-rich Task Suite and PETS-Net for In-hand Manipulation

The task suite is organized as a unified progression from simpler contact-rich behaviours to highly dynamic dexterous manipulation. It includes established single-arm tactile tasks such as object pushing, edge following, and surface following [2], [3]; bimanual tactile tasks such as bi-pushing [9]; shear-sensitive tactile tracking tasks such as dual-arm object following [10]; and dexterous in-hand manipulation tasks such as gravity-invariant object rotation [11]. By drawing from and extending prior tactile learning benchmarks, TactileLab unifies non-prehensile interaction, dual-arm coordination, precision contact alignment, and multi-finger dexterity within a single framework.

While the broader task suite demonstrates the generality of the tactile simulation framework, we focus PETS-Net on the most challenging dexterous in-hand manipulation setting. This task requires long-horizon reasoning over multi-finger contacts, proprioceptive histories, previous actions, and high-dimensional tactile observations. PETS-Net is introduced as a Positional-Encoding Temporal-Spatial Network for this setting, where it serves as the student adaptation module in a teacher–student framework. It applies positional encoding to joint-angle and action histories, extracts spatial features from contact-depth and tactile-flow observations, and temporally aggregates these features to estimate the privileged latent representation. This design allows the in-hand manipulation policy to exploit dense tactile information over time while remaining deployable without privileged simulation states.

D. Real2Sim Tactile Transfer and Sim2Real Policy Deployment

To bridge the gap between simulated and real tactile observations, we develop a Real2Sim tactile transfer pipeline that maps real camera-based tactile images into the simulated tactile modalities used during policy training. The transfer module predicts simulation-style contact-depth images and tactile-flow fields from real tactile observations, reducing the observation mismatch between simulation and the real robot. In particular, tactile-flow transfer uses consecutive real tactile images as input and predicts the corresponding simulated tactile-flow representation, allowing dynamic tangential contact cues observed by real vision-based tactile sensors to be converted into the modality used by TactileLab policies.

This transfer formulation is naturally suited to compliant vision-based tactile sensors, where contact deformation and shear can produce observable image-space changes. Instead of requiring high-fidelity soft-body rendering during reinforcement learning, TactileLab trains policies with efficient simulated tactile observations and later aligns real tactile images to these modalities through supervised Real2Sim translation. Combined with GPU-parallelized simulation and scalable tactile data collection, this provides an end-to-end workflow from simulated policy training to real-world tactile deployment.

III. Conclusion

TactileLab presents a unified framework for efficient shear-sensitive tactile robot learning. By combining GPU-parallelized tactile simulation, simulated tactile flow, multi-modal tactile policy learning, Real2Sim tactile transfer, and a broad task suite across multiple embodiments, the framework supports scalable learning for contact-rich manipulation from simulation to real deployment.

References

[1] Xuyang Zhang, Jiaqi Jiang, Zhuo Chen, Yongqiang Zhao, Tianqi Yang, Daniel Fernandes Gomes, Jianan Wang, and Shan Luo. Simtac: A physics-based simulator for vision-based tactile sensing with biomorphic structures. *Cyborg and Bionic Systems*, 7:0510, 2026.

[2] Alex Church, John Lloyd, Raia Hadsell, and Nathan F Lepora. Tactile sim-to-real policy transfer via real-to-sim image translation. In *Conference on Robot Learning*, pages 1645–1654, 2022.

[3] Yijiong Lin, John Lloyd, Alex Church, and Nathan F Lepora. Tactile gym 2.0: Sim-to-real deep reinforcement learning for comparing low-cost high-resolution robot touch. *IEEE Robotics and Automation Letters*, 7(4):10754–10761, 2022.

[4] Daniel Fernandes Gomes, Paolo Paoletti, and Shan Luo. Generation of gelsight tactile images for sim2real learning. *IEEE Robotics and Automation Letters*, 6(2):4177–4184, 2021.

[5] Shaoxiong Wang, Mike Lambeta, Po-Wei Chou, and Roberto Calandra. Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation Letters*, 7(2):3930–3937, 2022.

[6] Quan Khanh Luu, Nhan Huu Nguyen, et al. Simulation, learning, and application of vision-based tactile sensing at large scale. *IEEE Transactions on Robotics*, 39(3):2003–2019, 2023.

[7] Iretiayo Akinola, Jie Xu, Jan Carius, Dieter Fox, and Yashraj Narang. Tacsl: A library for visuotactile sensor simulation and learning. *IEEE Transactions on Robotics*, 2025.

[8] Mayank Mittal, Pascal Roth, James Tigue, Antoine Richard, Octi Zhang, Peter Du, Antonio Serrano-Muñoz, Xinjie Yao, René Zurbrügg, Nikita Rudin, Lukasz Wawrzyniak, Milad Rakhsha, Alain Denzler, Eric Heiden, Ales Borovicka, Ossama Ahmed, Iretiayo Akinola, Abrar Anwar, Mark T. Carlson, Ji Yuan Feng, Animesh Garg, Renato Gasoto, Lionel Gulich, Yijie Guo, M. Gussert, Alex Hansen, Mihir Kulkarni, Chenran Li, Wei Liu, Viktor Makoviychuk, Grzegorz Malczyk, Hammad Mazhar, Masoud Moghani, Adithyavairavan Murali, Michael Noseworthy, Alexander Poddubny, Nathan Ratliff, Welf Rehberg, Clemens Schwarke, Ritvik Singh, James Latham Smith, Bingjie Tang, Ruchik Thaker, Matthew Trepte, Karl Van Wyk, Fangzhou Yu, Alex Millane, Vikram Ramasamy, Remo Steiner, Sangeeta Subramanian, Clemens Volk, CY Chen, Neel Jawale, Ashwin Varghese Kuruttukulam, Michael A. Lin, Ajay Mandlekar, Karsten Patzwaldt, John Welsh, Huihua Zhao, Fatima Anes, Jean-Francois Lafleche, Nicolas Moënne-Loccoz, Soowan Park, Rob Stepinski, Dirk Van Gelder, Chris Amevor, Jan Carius, Jumyung Chang, Anka He Chen, Pablo de Heras Ciechomski, Gilles Daviet, Mohammad Mohajerani, Julia von Muralt, Viktor Reutskyy, Michael Sauter, Simon Schirm, Eric L. Shi, Pierre Terdiman, Kenny Vilella, Tobias Widmer, Gordon Yeoman, Tiffany Chen, Sergey Grizan, Cathy Li, Lotus Li, Connor Smith, Rafael Wiltz, Kostas Alexis, Yan Chang, David Chu, Linxi "Jim" Fan, Farbod Farshidian, Ankur Handa, Spencer Huang, Marco Hutter, Yashraj Narang, Soha Pouya, Shiwei Sheng, Yuke Zhu, Miles Macklin, Adam Moravanszky, Philipp Reist, Yunrong Guo, David Hoeller, and Gavriel State. Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025.

[9] Yijiong Lin, Alex Church, Max Yang, Haoran Li, John Lloyd, Dandan Zhang, and Nathan F. Lepora. Bi-touch: Bimanual tactile manipulation with sim-to-real deep reinforcement learning. *IEEE Robotics and Automation Letters*, 8(9):5472–5479, 2023.

[10] John Lloyd and Nathan F Lepora. Pose-and-shear-based tactile servoing. *The International Journal of Robotics Research*, 43(7):1024–1055, 2024.

[11] Max Yang, Chenghua Lu, Alex Church, Yijiong Lin, Chris Ford, Haoran Li, Efi Psomopoulou, David AW Barton, and Nathan F Lepora. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. *arXiv preprint arXiv:2405.07391*, 2024.